arXiv:1402.1810v1 [cs.DS] 8 Feb 2014

A Natural Generalization of Bounded Tree-Width and Bounded Clique-Width

Martin Fürer*

Department of Computer Science and Engineering
Pennsylvania State University
University Park, PA 16802, USA
furer@cse.psu.edu
http://cse.psu.edu/~furer

Abstract. We investigate a new width parameter, the fusion-width of a graph. It is a natural generalization of the tree-width, yet strong enough that not only graphs of bounded tree-width, but also graphs of bounded clique-width, trivially have bounded fusion-width. In particular, there is no exponential growth between tree-width and fusion-width, as is the case between tree-width and clique-width. The new parameter gives a good intuition about the relationship between tree-width and clique-width.

Keywords: tree-width, clique-width, fusion-width, FPT, XP

1 Introduction

Tree-width is a very natural concept. In an intuitive direct way, it measures how similar a graph is to a tree. Many graph problems are not only easy for trees, but also for other tree-like graphs. Indeed there is a huge number of efficient algorithms for graphs of bounded tree-width.

While graphs of bounded tree-width are sparse, there are some dense graphs, like the complete graph K_n or the complete bipartite graph K_{nn} for which most computational problems have a trivial solution. Graphs of bounded clique-width are intended to cover classes of graphs for which many problems have efficient solutions, even though they contain many dense graphs.

Not unlike tree-width, the concept of clique-width [11] is based on a type of graph decompositions [7] chosen to allow fast algorithms for large classes of graphs. As the name clique-width indicates, this measure is designed to ensure that complete graphs have a very small width. But neither does clique-width measure a closeness to a clique in a natural way, nor is there an intuitive width (as in tree-width) involved in the definition of this concept. Even though the definition of clique-width is fairly simple, it is harder to obtain an intuition for the classes of graphs with small clique-width.

* Research supported in part by NSF Grant CCF-0964655 and CCF-1320814.

Bounded clique-width is an extension of bounded tree-width in the sense that every class of bounded tree-width is also a class of bounded clique-width [11,6]. But the worst case clique-width of graphs of tree-width k has been upper [11,6] and lower [6] bounded by an exponential in k. The fact that this containment result is a difficult theorem suggests that the extension from tree-width to clique-width might not be very natural.

Equivalent to clique-width up to a factor of 2 is the notion of *k-NLC (node label controlled) graphs*. Partial k-trees have been shown to be $(2^{k+1}-1)$-NLC trees [21]. k-NLC trees are a sub-class of the k-NLC graphs.

In contrast to clique-width (and the width measure produced by NLC trees), we propose a natural generalization of tree-width, which simultaneously generalizes clique-width. Furthermore, containment in the new class is obtained basically without increasing the parameter in both cases. We call the new measure *fusion-width*. We initially choose the name multi-tree-width to emphasize that it is a very natural extension of tree-width, by which it is motivated (even though it is much more powerful). We follow the strong suggestion of two referees to name it differently. The main difference is that while tree-width deals with single vertices or pairs of vertices at a time, fusion-width deals with multiple vertices (with the same label) or pairs of sets of vertices at a time.

We show that the clique-width grows at most exponentially in the fusion-width, implying that all classes of graphs with bounded fusion-width also have bounded clique-width. Furthermore, for some classes of graphs, there is really such an exponential growth.

Other width parameters generalizing both tree-width and clique-width without blowing up the parameter are rank-width [18] and boolean width [5]. The rank-width has the additional nice property of being computable in FPT [15]. For an overview of width parameters, see [16]. There are infinite classes of graphs where the clique-width is exponentially bigger than the boolean width.

The fusion-width has the additional property of being easy to work with and of being the most natural generalization of tree-width and clique-width. This is an essential strength of the new notion of fusion-width.

2 Definitions

For the definitions of *FPT (fixed parameter tractable)*, *XP (fixed parameter polynomial time)*, and *tree decomposition*, see e.g. [12].

Definition 1. *The* tree-width *$tw(G)$ [19] of a graph G is the smallest integer k, such that G has a tree decomposition with largest bag size $k+1$.*

It is NP-complete to decide whether the tree-width of a graph is at most k (if k is part of the input) [1]. For every fixed k, there is a linear time algorithm deciding whether the tree-width is at most k, and if that is the case, producing a corresponding tree decomposition [2]. For arbitrary k, this task can still be approximated. A tree decomposition of width $O(k \log n)$ [4] and even $5k$ [3] can be found in polynomial time.

Closely related to tree-width is the notion of *branch-width* [20].

Definition 2. *A* k-expression *is an expression formed from the atoms* $i(v)$*, the two unary operations* $\eta_{i,j}$ *and* $\rho_{i\to j}$*, and one binary operation* $\oplus$ *as follows.*

- $i(v)$ *creates a vertex* v *with label* i*, where* i *is from the set* $\{1,\ldots,k\}$*.*
- $\eta_{i,j}$ *creates an edge between every vertex with label* i *and every vertex with label* j *for* $i \neq j$*.*
- $\rho_{i\to j}$ *changes all labels* i *to* j*.*
- $\oplus$ *does a disjoint union of the generated labeled graphs.*

Finally, the generated graph *is obtained by deleting the labels.*

Definition 3. *The* clique-width $cw(G)$ *of a graph is the smallest* k *such that the graph can be defined by a* k*-expression [7,11].*

Computing the clique-width is NP-hard [13]. Thus, one usually assumes that a graph is given together with a k-expression. For constant k, the clique-width can be approximated by a constant factor in polynomial time [18], in fact, this factor can be made smaller than 3 [17].

3 The Fusion-Width

We define a new width measure $\mathrm{fw}(G)$ (fusion-width of G) with the properties

$$\mathrm{fw}(G) \leq \mathrm{tw}(G) + 2 \text{ and } \mathrm{fw}(G) \leq \mathrm{cw}(G).$$

We want these containments to be obvious. Still, we would like all tasks known to be solvable in polynomial time for graphs of bounded clique-width (and therefore of bounded tree-width) to be solvable in polynomial time also for graphs of bounded fusion-width.

These inequalities immediately imply that the class of graphs of bounded clique-width is contained in the class of graphs of bounded fusion-width. The definition of fusion-width is obtained as a simple extension of the definition of clique-width by a new operation θ_i, merging all vertices with label i.

Definition 4. *A* k-fusion-tree expression *is an expression formed from the atoms* $i\langle m\rangle$*, the three unary operations* $\eta_{i,j}$*,* $\rho_{i\to j}$*, and* θ_i*, and one binary operation* $\oplus$ *as follows.*

- $i\langle m\rangle$ *creates a graph consisting of* m *isolated vertices labeled* i*, where* i *is from the set* $\{1,\ldots,k\}$*.*
- $\eta_{i,j}$ *creates an edge between every vertex with label* i *and every vertex with label* j *for* $i \neq j$*.*
- $\rho_{i\to j}$ *changes all labels* i *to* j*.*
- θ_i *merges all vertices labeled* i *into one vertex. The new vertex is labeled* i *and is adjacent to every vertex not labeled* i *to which some vertex labeled* i *was adjacent before the operation.*

– $\oplus$ *does a disjoint union.*

Finally, the generated graph *is obtained by deleting all the labels.*

It might be more natural not to require $i \neq j$ for $\eta_{i,j}$. This is insignificant, but would have the nice effect of giving any clique a fusion-width of 1 instead of 2. Thus the simplest graphs in this measure would just be the collections of disjoint cliques. Nevertheless, we stick to the traditional $\eta_{i,j}$ operation.

Definition 5. *The* fusion-width *fw*(G) *of a graph* G *is the smallest* k *such that there is a* k*-fusion-tree expression generating* G*.*

The operation $i\langle m\rangle$ is introduced for convenience and to emphasize the possibility to create huge collections of identical vertices and huge bipartite graphs (together with $\eta_{i,j}$). We always want to emphasize the difference between fusion-width and tree-width here. Otherwise, except for the θ_i operation, we have just the clique-width operations, and the slightly more efficient version of vertex creation.

Compared to clique-width, the definition of fusion-width contains the completely new operation θ_i. It is introduced to directly mimic the tree-width construction. An immediate consequence is that bounded tree-width graphs have also bounded fusion-width, without a difficult proof and without an exponential blow-up in the width parameter. This is in sharp contrast to the relationship between bonded tree-width and bounded clique-width.

Thanks to a previous referee, we know that Courcelle and Makowsky [8] have already defined the parameter $\mathrm{fw}(G)$ when they showed that labelled graphs of bounded clique-width are closed their fusion operator. They call the parameter $\mathrm{cwd}'(G)$, viewing it as an alternative notion of clique-width (justified by bounded clique-width being equivalent to bounded $\mathrm{cwd}'(G)$). They don't propose $\mathrm{cwd}'(G)$ to be used as a new width measure nor do they relate $\mathrm{cwd}'(G)$ to tree-width.

Theorem 1. *Graphs with clique-width* k *have fusion-width at most* k*. Furthermore, the number of operations does not increase from the associated* k*-expression to the associated* k*-fusion-tree expression.*

Proof. This is trivial, because all the operations of k-expressions are also operations of the k-fusion-tree expressions (with the obvious variation of replacing $i(v)$ by $i\langle 1\rangle$). □

The width parameter increases by at most 2 from tree-width to fusion-width. This tiny increase has two causes. An increase of 1 is just due to the somewhat artificial push-down by 1 in the definition of tree-width. We don't use it for fusion-width, because we want to align the measure with clique-width. The other increase by 1 is needed to have an extra label for the vertices that have already received all their incident edges.

Theorem 2. *Graphs of tree-width* k *have fusion-width at most* $k + 2$*.*

Proof. We start with a tree decomposition with bag size $k+1$, and transform it into a $k+2$-fusion-tree expression in a bottom-up way. One special label is reserved for vertices that have already been handled, i.e., all their incident edges have been produced. Here we refer to the $k+1$ other labels as regular labels.

In each bag, we use different labels for different vertices. Thus, when handling a bag, it is trivial to introduce all edges present in that bag by the corresponding $\eta_{i,j}$ operations.

Choosing such a labeling is easy to do top down. We select an arbitrary node as the root of the decomposition tree, and assign distinct labels to the vertices in its bag. Before we assign labels to vertices in the bag of a node j not appearing in the bag of the parent node, we assume that every vertex appearing in the bag of j as well as in the bag of its parent node of the decomposition tree has already received its label. If there are still vertices in the bag of node j without a label, then there are enough unused regular labels for these vertices, because we have $k+1$ regular labels and at most $k+1$ vertices in the bag of j.

The only slightly tricky part of the fusion-tree expression is the handling of the fact that the same vertices occur in the bags of both children of a node in the tree decomposition. This is handled by using distinct vertices with the same label. When handling a node in the decomposition tree, such that more than one of its children contain the same vertex v in their bags, a merge operation θ_i is issued with i being the common label of all occurrences of v in the subtrees. □

Naturally, the following corollary is an immediate consequence.

Corollary 1. *If a problem can be solved in time $T(n,k)$ for graphs with n vertices and fusion-width k, then it can be solved in time $T(n,k+2)$ for graphs with n vertices and tree-width k.*

This corollary should be compared with the corresponding important result for clique-width.

If a problem can be solved in time $T(n,k)$ for graphs with n vertices and clique-width k, then it can be solved in time $T(n,3\cdot 2^{k-1})$ for graphs with n vertices and tree-width k [6].

Instead of using this result, one would rather apply the corollary with its much stronger conclusion, provided that the premises are comparable.

We believe that whenever there is a nice argument that a natural problem can be solved for graphs of bounded clique-width, then we are able to nicely handle the operations $\eta_{i,j}$, $\rho_{i\to j}$, and $\oplus$. Usually, we would then also have a nice argument that the problem could be solved for graphs of bounded tree-width, and we could nicely handle the operation θ_i. In such a situation, we would be able to handle all the fusion-width operations nicely, and therefore would also have an elegant algorithm whose running time should not be too bad as a function of the fusion-width.

The allowance of merging vertices with the θ_i operation might cause two concerns. First, it is more powerful than necessary for our results. It would be sufficient to restrict it to sets of vertices of size 2. Nevertheless, we opted for the more flexible notion, because it does not cause any problems. A second concern

looks more important. As the construction of graphs allows them to grow and shrink, it is reasonable to ask whether there are graphs of bounded fusion-width requiring super-polynomial size k-fusion-tree expressions. This is not the case

Proposition 1. *Every graph of fusion-width k has a k-fusion-tree expression of size $O(|V| + |E|)$.*

Proof. Whenever some vertices are merged due to a θ_i operation, it is possible that some edges are merged too. Assume, there is a vertex v that has been used to create some edge set E_v with some $\eta_{i,j}$ operations. Further assume that when v is merged with some set of vertices, every edge of E_v is merged with at least one other edge. Then we obtain the same graph by omitting the creation of vertex v. In other words, every vertex ever created is either useless, or it is an isolated vertex in the resulting graph $G = (V, E)$, or it is responsible for at least one edge. Now, assume no useless vertices (that have no effect and disappear in a later merge operation) are ever created. Then the number of vertices ever created is at most $|V|+|E|$. Furthermore, it is obvious, that without unnecessary label change operations $\rho_{i\to j}$, the graph G has a k-fusion-tree expression of size $O(|V| + |E|)$. □

4 Illustration with the Independent Set Polynomial

Naturally, we know that finding a maximum independent set is possible in polynomial time for graphs of bounded clique-width [9]. In fact the far reaching meta-theorem of Courcelle et al. [9] shows that this result is not just valid for the maximum independent set problem, but for every problem expressible in monadic second order logic with quantification only over sets of vertices (not edges). Furthermore, the resulting algorithm shows the problem to be in FPT with the clique-width as the parameter.

Here, we look at a more difficult problem. Instead of just finding the size of a maximum independent set for graphs of bounded clique-width, we count the number of independent sets of all sizes. We present a fixed parameter polynomial time algorithm for this counting problem. We refer to [10,14] for more discussions of the fixed parameter tractability of counting problems.

Let $[k] = \{1, \ldots, k\}$ be the set of vertex labels. We define the $[k]$-labeled independent set polynomial of a $[k]$-labeled graph G by

$$P(x, x_1, \ldots, x_k) = \sum_{i=1}^{n} \sum_{(n_1,\ldots,n_k)\in\{0,1\}^k} a_{i;n_1,\ldots,n_k}\, x^i \prod_{j=1}^{k} x_j^{n_j}$$

where $n_j \in \{0, 1\}$ and $a_{i;n_1,\ldots,n_k}$ is the number of independent sets of size i in G which contain some vertices with label j if and only if $n_j = 1$.

We define the independent set polynomial of a graph G by

$$I(x) = \sum_{i=1}^{n} a_i x^i$$

where a_i is the number of independent sets of size i in G.

Then the independent set polynomial $I(x)$ can immediately be expressed by the $[k]$-labeled independent set polynomial $P(x, x_1, \ldots, x_k)$.

Proposition 2. *The independent set polynomial $I(x)$ of a $[k]$-labeled graph G is*

$$I(x) = \sum_{(n_1,\ldots,n_k)\in\{0,1\}^k} P(x, 1, \ldots, 1) = \sum_{i=1}^{n} \sum_{(n_1,\ldots,n_k)\in\{0,1\}^k} a_{i,n_1,\ldots,n_k} x^i.$$

Theorem 3. *Given a graph G with n vertices and bounded fusion-width k, and a polynomial size k-fusion-tree expression generating G, the independent set polynomial $I(x)$ of G can be computed in FPT, i.e., in time $f(k)n^{O(1)}$ for some function f.*

Proof. We have to show how to compute the $[k]$-labeled independent set polynomial $P(x, x_1, \ldots, x_k)$ of a $[k]$-labeled graph G. We compute it recursively bottom-up for the given k-fusion-tree expression. For the edgeless graph with m vertices and label i generated by $i\langle m\rangle$, the $[k]$-labeled independent set polynomial is

$$x + \sum_{j=1}^{m} \binom{m}{j} x^j x_i.$$

This polynomial is computable in time polynomial in n, because w.l.o.g, we can assume $m \leq n$. Otherwise, some set of vertices would be constructed together (by the $i\langle m\rangle$ operation) and destroyed together (with the same merge operation θ_i). Such a redundancy can easily be removed in a preprocessing phase.

In the following, assume for some $[k]$-labeled graph H, the $[k]$-labeled independent set polynomial is $\tilde{P}(x, x_1, \ldots, x_k)$.

Let G be obtained from H by the operation $\eta_{i,j}$. Then the $[k]$-labeled independent set polynomial $P(x, x_1, \ldots, x_k)$ of G is obtained from the $[k]$-labeled independent set polynomial $\tilde{P}(x, x_1, \ldots, x_k)$ of H by deleting all monomials that are multiples of $x_i x_j$. These monomials count sets that are no longer independent after inserting all the edges between vertices labeled i and j.

W.l.o.g., we can assume that before a merge operation θ_i is done, there are only 2 vertices labeled i. This assumption is allowed for two reasons.

1. If later a θ_i operation is done, then every $i\langle m\rangle$ operation can be replaced by an $i\langle 1\rangle$ operation without changing the graphs obtained after the θ_i operation. Creating many equivalent vertices (with the same neighbors) and merging them later has the same effect as creating just one vertex.
2. Every θ_i operation can be replaced by a collection of θ_i operations done immediately after a disjoint union $\oplus$ or a relabeling operation $\rho_{i\to j}$ has created a graph with two vertices labeled i.

We describe the θ_i operation not in isolation, but only immediately after a disjoint union $\oplus$ or a relabeling operation $\rho_{i\to j}$.

We now describe how to obtain the polynomial $P(x, x_1, \ldots, x_k)$ of G from the polynomials $P_r(x, x_1, \ldots, x_k)$ of H_r $(r = 1, 2)$, where G is obtained from H_1 and H_2 by the operation $\theta_{i_1} \ldots \theta_{i_\ell}(H_1 \oplus H_2)$. For ease of notation, assume that $\{x_1, \ldots, x_\ell\} = \{x_{i_1}, \ldots, x_{i_\ell}\}$.

- Form the product $P_1(x, x_1, \ldots, x_k) \cdot P_2(x, x_1, \ldots, x_k)$.
- Delete all monomials where some x_j with $1 \leq j \leq \ell$ appears with an exponent of 1.
  If u and v merge into vertex w, then we either want both u and v in the independent set (to account for an independent set containing w), or neither (to account for an independent set not containing w).
- Replace x_j^2 by x_j/x for $1 \leq j \leq \ell$.
  Division by x compensates the double count of a vertex in the independent set (counting u and v for w).
- Replace x_j^2 by x_j for $\ell + 1 \leq j \leq k$.
  If label j is not merged, then x_j just indicates whether there are any vertices labeled j in the independent set.

The case of a simple disjoint union ($G = H_1 \oplus H_2$) is the special case $\ell = 0$ of the just described situation. Here, we just compute the product $P_1(x, x_1, \ldots, x_k) \cdot P_2(x, x_1, \ldots, x_k)$ and delete all monomials x_j (for $1 \leq j \leq k$) appearing with an exponent 1 to obtain $P_r(x, x_1, \ldots, x_k)$.

We now consider the relabeling operation $\rho_{i \to j}$. First assume, there will be no succeeding θ_j operation. Let G be obtained from H by the operation $\rho_{i \to j}$. Then the $[k]$-labeled independent set polynomial $P(x, x_1, \ldots, x_k)$ of G is obtained from $\tilde{P}(x, x_1, \ldots, x_k)$ by substituting x_j for x_i and then replacing x_j^2 by x_j.

If G is obtained from H by the operation $\theta_j \rho_{i \to j}$, then we assume that in H there is just one vertex labeled i and one vertex labeled j. In this case, we proceed as follows, with the same reasoning as for the disjoint union $\oplus$ operation.

- In the given polynomial $\tilde{P}(x, x_1, \ldots, x_k)$ of H, substitute x_j for x_i.
- Delete all monomials where x_j appears with an exponent of 1.
- Replace x_j^2 by x_j.

To prove the polynomial time claim, it is important to notice that all polynomials have at most $k + 1$ variables, and all monomials are linear in each of their variables. Thus there are at most 2^{k+1} monomials. The number of arithmetic operations is $O(k^2 n)$, as every efficient k-fusion-tree expression has at most $O(k^2)$ unary operations in a row. Thus for k a constant, the time is at most $O(k^2 2^{2k} n)$ if a trivial polynomial multiplication algorithm is used. With fast polynomial multiplication, based on fast Fourier transform, the time goes down to $O(k^3 2^k n)$. As the total number of independent sets is at most 2^n, it is sufficient to do the computations with numbers of length $O(n)$. Thus, each arithmetic operation requires even with school multiplication only quadratic time. □

Note that we do not claim that this was a difficult theorem. To the contrary, the point was to illustrate that a fast algorithm for a typical problem like computing the independent set polynomial restricted to bounded clique-width

can be extended to a fast algorithm for this problem for graphs with the same fusion-width, i.e., for a much larger class of graphs.

5 Relations between Tree-Width, Clique-Width and Fusion-Width

We have $\mathrm{fw}(G) \leq \mathrm{tw}G + 2$ by Theorem 2. The following inequality is trivial, as k-fusion-tree expressions are strictly more powerful than k-expressions.

Proposition 3. *[8] $fw(G) \leq cw(G)$.*

The following main result immediately implies that the graphs of bounded clique-width are exactly the graphs of bounded fusion-width. In fact this implication has already been shown by Courcelle and Makowsky [8], as they prove the existence of a function f with $\mathrm{cw}(G) \leq f(\mathrm{fw}(G))$. We still present our direct proof, because we get a much stronger result, and also because the logic framework of [8] might not be so widely accessible. Many of the proof ideas are from the corresponding Theorem of Corneil and Rotics [6], relating tree-width to clique-width.

Theorem 4. *Graphs with fusion-width $fw(G) = k$ have clique-width $cw(G)$ at most $k2^k$.*

Proof. We assume, we are given a k-fusion-tree expression E describing a graph G, and we want to construct a $k2^k$-expression describing the same graph G.

We have to focus on the operation θ_i merging all vertices labeled i into one vertex. This is the only operation that has to be eliminated, because it is allowed in k-fusion-tree expressions determining the fusion-width, but not in k-expressions determining the clique-width.

We view the parse trees T of the k-fusion-tree expression. We want to transform it into a parse tree T' of a k-expression. The main idea is that if the vertices of some label i are involved in a merge operation, we focus on the highest location ℓ in T where such a merge occurs involving label $\hat{i}$, where $\hat{i}$ is either i or a label i has been changed to.

At the corresponding location ℓ' in T', we create the single vertex v to which the vertices labeled $\hat{i}$ have been merged by $\theta_{\hat{i}}$ using the operation $\hat{i}\langle 1\rangle$. This means that all the vertices $v_1, \ldots v_p$ which are finally merged into v are not available further down in the tree T'. Therefore all operations in T involving the labels of $v_1, \ldots v_p$ have to be delayed until the vertex creation operation at ℓ'.

Let L be the set of labels used in T. The new labels in T' are from $L \times \mathcal{P}(L)$, where $\mathcal{P}$ is the powerset of L. This way, every new label can retain its own (old) identity and in addition remember all the other old labels to which its vertices should actually be adjacent according to the edge constructing operations $\eta_{i,j}$ issued in the subtree of the current node in T.

Whenever a label changes due to a renaming operation $\rho_{i\to j}$, in T, then every occurrence of i (in the first or second component) of a label in T' is changed to j.

We say that a label $i = i_1$ is subject to a merge operation as label $\hat{i} = i_{q+1}$, if there is a sequence of label change operations $\rho_{i_p \to i_{p+1}}$ $(p = 1, 2, \ldots, q,\ q \geq 0)$ (i.e., there might be no label changes and $\hat{i} = i$), such that after these changes, label i_{q+1} is involved in a merge operation.

Now, any $\eta_{i,j}$ operation in T is handled as follows.

- If neither label i nor j are subject to a merge operation, then in the corresponding location (involving several nodes) of T', $\eta_{i',j'}$ is issued for all labels i' with first component i and labels j' with first component j.
- If label i is subject to a merge operation as label $\hat{i}$, and j is not subject to a merge operation, then the label i is added to the second component of j. This way nodes labeled j "remember" to create an edge to label $\hat{i}$ later. This edge is created immediately after the last merge of label $\hat{i}$.
- If label i is subject to a last merge operation as label $\hat{i}$ after j has been subject to its last merge operation, then the label i is added to the second component of j. This way nodes labeled j "remember" to create an edge to label $\hat{i}$ later. This edge is created immediately after the last merge of label $\hat{i}$.

□

Corollary 2. *A class of graphs is of bounded fusion-width if and only if it is a class of bounded clique-width.*

Proof. This follows immediately from Proposition 3 and Theorem 4. □

Corollary 3. *If a problem can be solved in time $O(f(k)n^c)$ for graphs of fusion-width at most k, then it can be solved in time $O(f(k+2)n^c)$ for graphs of tree-width at most k.*

This immediate corollary compares favorably with the following fact. If a problem can be solved in time $O(f(k)n^c)$ for graphs of clique-width at most k, then it can be solved in time $O(2^{f(k)}n^c)$ for graphs of tree-width at most k.

This statement cannot be much improved, as clique-width can be exponentially larger than fusion-width.

Theorem 5. *[6] For any k, there is a graph G with $tw(G) = k$ and $cw(G) \geq 2^{\lfloor k/2 \rfloor - 1}$.*

Corollary 4. *For any k, there is a graph G with fusion-width $fw(G) = k$ and clique-width $cw(G) \geq 2^{\lfloor k/2 \rfloor - 2}$.*

Proof. This follows from Theorem 2 and Theorem 5 by picking a graph with the properties of Theorem 5 and noticing that by Theorem 2 its fusion-width is at most $k + 2$. Then $2^{\lfloor (k-2)/2 \rfloor - 1} = 2^{\lfloor k/2 \rfloor - 2}$ produces the result. □

Corollary 4 shows that our example of the independent set polynomial proves the fusion-width to be a powerful notion. There are graphs with fusion-width $\mathrm{fw}(G) = k$ and clique-width $\mathrm{cw}(G) \geq 2^{\lfloor k/4 \rfloor - 2}$. If for such a graph, we have a k-fusion-tree expression, then we can compute its independent set polynomial

in time $O(k^3 2^k n)$ by the method of Theorem 3. Using just the fact that the clique-width $\mathrm{cw}(G) \geq 2^{\lfloor k/4 \rfloor - 2}$, we would be unlikely to find a better algorithm based on clique-width. Thus we would only have an an algorithm that is doubly exponential in k for computing the independent set polynomial of these graphs.

We believe that the independent set polynomial is not an isolated instance. It has merely been used to illustrate the convenience and power of the fusiion-with parameter. Many other examples could be used instead.

6 Conclusion

We have introduced a new width measure, the fusion-width. Its purpose is two-fold. It provides a tool for handling generally difficult problems for a large class of graphs. It also sheds a light on the essence of the generalization from bounded tree-width to bounded clique width. It is the ability at each stage of the construction not only to add edges between a limited number of vertices, but to add complete bipartite graphs between a limited number of sets of vertices.

7 Open Problems

What is the complexity of determining the fusion-width of a graph? Is it in XP (fixed parameter polynomial time) or even in FPT (fixed parameter tractable)? How well can fusion-width be approximated?

Find interesting classes of graphs with a large clique-width and a small fusion-width.

What is the relationship between fusion-width, rank-width, and boolean width?

8 Acknowlegement

The help of previous anonymous reviewers has improved this paper significantly.

References

1. Arnborg, S., Corneil, D.G., Proskurowski, A.: Complexity of finding embeddings in a k-tree. SIAM Journal of Alg. and Discrete Methods 8, 277–284 (1987)
2. Bodlaender, H.L.: A linear-time algorithm for finding tree-decompositions of small treewidth. SIAM J. Comput. 25(6), 1305–1317 (1996)
3. Bodlaender, H.L., Drange, P.G., Dregi, M.S., Fomin, F.V., Lokshtanov, D., Pilipczuk, M.: An $O(c^k n)$ 5-approximation algorithm for treewidth. In: Proc. 54th FOCS 2013. pp. 499–508. IEEE (2013)
4. Bodlaender, H.L., Gilbert, J.R., Hafsteinsson, H., Kloks, T.: Approximating treewidth, pathwidth, frontsize, and shortest elimination tree. J. Algorithms 18(2), 238–255 (1995)
5. Bui-Xuan, B.M., Telle, J.A., Vatshelle, M.: Boolean-width of graphs. Theor. Comput. Sci. 412(39), 5187–5204 (2011)

6. Corneil, D.G., Rotics, U.: On the relationship between clique-width and treewidth. SIAM J. Comput. 34(4), 825–847 (2005)
7. Courcelle, B., Engelfriet, J., Rozenberg, G.: Handle-rewriting hypergraph grammars. J. Comput. Syst. Sci. 46(2), 218–270 (1993)
8. Courcelle, B., Makowsky, J.A.: Fusion in relational structures and the verification of monadic second-order properties. Mathematical Structures in Computer Science 12(2), 203–235 (2002)
9. Courcelle, B., Makowsky, J.A., Rotics, U.: Linear time solvable optimization problems on graphs of bounded clique-width. Theory Comput. Syst. 33(2), 125–150 (2000)
10. Courcelle, B., Makowsky, J.A., Rotics, U.: On the fixed parameter complexity of graph enumeration problems definable in monadic second-order logic. Discrete Applied Mathematics 108(1-2), 23–52 (2001)
11. Courcelle, B., Olariu, S.: Upper bounds to the clique width of graphs. Discrete Applied Mathematics 101(1-3), 77–114 (2000)
12. Downey, R., Fellows, M.R.: Parameterized complexity. Monographs in computer science., Springer, New York (1999), r.G. Downey, M.R. Fellows. New Zealand authors. Includes bibliographical references (p. [489]-516) and index.
13. Fellows, M.R., Rosamond, F.A., Rotics, U., Szeider, S.: Clique-width minimization is np-hard. In: Kleinberg, J.M. (ed.) STOC. pp. 354–362. ACM (2006)
14. Fischer, E., Makowsky, J.A., Ravve, E.V.: Counting truth assignments of formulas of bounded tree-width or clique-width. Discrete Applied Mathematics 156(4), 511–529 (2008)
15. Hlinený, P., Oum, S.: Finding branch-decompositions and rank-decompositions. In: Arge, L., Hoffmann, M., Welzl, E. (eds.) ESA. Lecture Notes in Computer Science, vol. 4698, pp. 163–174. Springer (2007)
16. Hlinený, P., Oum, S., Seese, D., Gottlob, G.: Width parameters beyond tree-width and their applications. Comput. J. 51(3), 326–362 (2008)
17. Oum, S.: Approximating rank-width and clique-width quickly. ACM Trans. Algorithms 5(1), 10:1–10:20 (Dec 2008), http://doi.acm.org/10.1145/1435375.1435385
18. Oum, S., Seymour, P.D.: Approximating clique-width and branch-width. J. Comb. Theory, Ser. B 96(4), 514–528 (2006)
19. Robertson, N., Seymour, P.D.: Graph minors. III. Planar tree-width. J. Comb. Theory, Ser. B 36(1), 49–64 (1984)
20. Robertson, N., Seymour, P.D.: Graph minors. X. Obstructions to tree-decomposition. J. Comb. Theory, Ser. B 52(2), 153–190 (1991)
21. Wanke, E.: k-NLC graphs and polynomial algorithms. Discrete Applied Mathematics 54(2-3), 251–266 (1994)